

*AERONAUTICAL SYSTEMS CENTER
MAJOR SHARED RESOURCE CENTER*



*SGI ORIGIN 3900
USER'S GUIDE*

April 06, 2004

Release Notes

Issued April 09, 2003

Internal Draft

Issued May 05, 2003

Internal Draft. Clarified and verified system information, corrected spelling mistakes, etc.

Issued May 15, 2003

Internal Draft/Pioneer Release. Corrected inconsistencies, spelling errors and incorrect information.

Issued July 10, 2003

Production release. Please report all errors, comments and feedback to msrchelp@asc.hpc.mil.

Issued July 22, 2003

Corrected errors in the parallel compile examples.

Issued September 23, 2003

Updated LSF script examples, updated information about CPH, added information about Jlimits.

Issued April 06, 2004

Updated links, clarified wording on /workspace structure, misc wording changes for clarity.

Table of Contents

1. Introduction.....	1
1.1 Assumed Background of the Reader.....	1
1.2 Hardware Overview	1
1.3 Accessing the SGI Origin 3900	1
1.4 ASC MSRC Connectivity	1
1.5 ASC MSRC Startup Files	2
1.6 The Archive Command.....	2
1.7 Jlimit	2
1.8 Additional Information	3
2. ASC MSRC SGI Origin 3900	4
2.1 Hardware.....	4
2.2 SGI Origin 3900 File System Overview	4
2.3 SGI IRIX Operating System	5
2.4 Available Software.....	5
2.5 Development	5
3. Program Development.....	7
3.1 Parallel Processing	7
3.2 FORTRAN Programming.....	8
3.3 C/C++ Programming.....	8
3.4 Libraries	8
4. Running Jobs on the SGI Origin 3900	10
4.1 Interactive Use	10
4.2 Batch Use	10
5. Customer Service	15
5.1 Customer Service Center	15
5.2 ASC MSRC Support	15
5.3 ASC MSRC Website	15
Appendix. SGI Origin 3900 Usage Hints.....	17
A.1 Runtime Considerations.....	17
A.2 Files and Filespace	17
A.3 Helpful SGI Origin 3900-related Websites.....	18

1. Introduction

This document provides an overview and introduction to using the SGI Origin 3900 system, located at the Aeronautical Systems Center (ASC) Major Shared Resource Center (MSRC). The ASC MSRC is located at Wright-Patterson Air Force Base, near Dayton, Ohio. This guide is intended to provide information so that customers who are familiar with the UNIX operating system can create and run their own programs, as well as use existing application software on the SGI Origin 3900 system.

1.1 Assumed Background of the Reader

It is assumed that the reader of this guide has a firm grasp of the concepts required to use the UNIX operating system and to program in either the C, C++, FORTRAN 77, FORTRAN 90, or FORTRAN 95 languages. It is also assumed that the reader has read the *ASC MSRC User's Guide*, which contains site specific information about the ASC MSRC. The *ASC MSRC User's Guide* is available from the ASC MSRC Service Center at 1-888-MSRC-ASC (1-888-677-2272), (937)255-0194, or DSN 785-0194. It is also available in postscript and pdf formats as described below in Section 1.7. UNIX training is periodically offered by the ASC MSRC.

1.2 Hardware Overview

The SGI Origin 3900 system is a Shared Memory system with 512 Central Processing Units (CPUs) per node. Each CPU is a 700 MegaHertz (MHz) R14000 MIPS processor and has a 32 kilobyte (KB) primary data cache, a 32 KB primary instruction cache, and 8 megabyte (MB) on-board cache.

The SGI Origin 3900 is front-ended by an SGI Origin 3800 for interactive use and batch submittal. The hardware specifications for this front-end system are slightly different than the batch systems. The front-end SGI Origin 3800 is a Shared Memory system with 12 CPUs. Each CPU is a 600 MHz R12000 MIPS processor and has a 32 KB primary data cache, a 32 KB primary instruction cache, and 8 MB on-board cache.

1.3 Accessing the SGI Origin 3900

The SGI Origin 3900 has a total of 2048 CPUs, divided into 4 nodes with 512 CPUs per node. The fully qualified hostname of the interactive system is *hpc11.asc.hpc.mil*.

Users are only permitted to login onto the interactive node of the system. The other nodes are for batch jobs only. Users submit their jobs on the front-end node and the batch system will automatically start their jobs on the other systems based on the load of the system. See the *ASC MSRC User's Guide* for instructions on accessing the machines.

1.4 ASC MSRC Connectivity

Since the SGI Origin 3900 is an integrated component of the ASC MSRC, user files are Network File System (NFS) mounted from the ASC MSRC High Availability File Server (HAFS) system to the SGI Origin 3900. When users log into a system, their home (\$HOME) directory (which will be the current directory immediately

after logging in) physically resides on the file server, but appears to be on the SGI Origin 3900. The ASC MSRC also supplies archival storage and visualization capabilities.

1.5 ASC MSRC Startup Files

All users are provided a `.cshrc` and `.login` file in their home directory. These files reference standard setup files, maintained by the site administrators in a central location, which set up a standard environment for all MSRC users. These files **should not** be modified.

To set up specific information for your SGI Origin 3900 session, such as environment variables, path information, terminal information, or command aliases, place the appropriate commands and information into files called `.personal.cshrc` and `.personal.login`. The standard startup files check your home directory for the existence of these files and executes them if found. Commands related to aliases, prompts, and some environment variables should go into `.personal.cshrc`, while commands related to the type of terminal you are using should go into `.personal.login`. See Section 3 of this guide for more details on the SGI Origin 3900 computing environment and the *ASC MSRC User's Guide* for more details on startup files.

1.6 The Archive Command

The **archive** command is a recently added tool to the ASC MSRC to help users with transferring files to and from \$ARC. The basic syntax for the archive command is:

```
archive get [getopts] file1 [file2 ...]  
archive put [putopts] file1 [file2 ...]
```

More information on the **archive** command can be found online via the *archive* man page (*man archive*).

1.7 Jlimit

The **jlimit** command displays and changes limits on resource usage for a job similar in many ways to the `limit` command. `Jlimit`, however, allows per user limits instead the same limits for all users. This allows resources to be modified dynamically, per user, without system downtime and without affecting other users on the system.

It can also assist support personnel in determining if your job is running into resource issues. To add `jlimit` debugging information to your job, place the following command before you call your executable and after you call your executable in your SGI LSF script:

```
jstat -l
```

This command will generate output similar to this:

LIMIT NAME	USAGE	HIGH USAGE	CURRENT LIMIT	MAX LIMIT
-----	-----	-----	-----	-----
cputime	0:01	0:01	unlimited	unlimited
datasize	608k	832k	unlimited	unlimited
files	23	63	100000	120000
vmemory	11936k	22560k	unlimited	unlimited
ressetsize	6848k	7632k	unlimited	unlimited
threads	0	0	unlimited	unlimited
processes	3	8	1024	2048
physmem	6848k	7632k	unlimited	unlimited

These are your jlimits, and can be changed using the *jlimit [limit_name] [value]* command and syntax. While jlimit allows values to be changed per user, there are limits set by the administrators that you cannot exceed. If you feel the maximum jlimit for the resource you need is not high enough, please contact the ASC MSRC Service Center asking that your limits be increased. Please note this is done on a case-by-case basis and may require justification.

1.8 Additional Information

Much of the information presented in this document is available online through the man pages and is accessible by typing:

man {command name}

when logged into any SGI Origin 3900 interactive node.

The *ASC MSRC User's Guide* and this document are all available in pdf and postscript format. They may be downloaded via the ASC MSRC website at

<http://www.asc.hpc.mil/customer/userdocs>

2. ASC MSRC SGI Origin 3900

This section details the hardware and software available on the SGI Origin 3900 and how they are currently configured.

2.1 Hardware

The SGI Origin 3900 has a total of 2048 CPUs. Each CPU is a 700 MHz MIPS R14000 processor with a peak speed of 1.2 GFLOPS, providing a total capacity of approximately 2.5 TFLOPs. Each CPU has a primary data cache of 32 KB, a primary instruction cache of 32 KB, and an on-board cache of 8 MB. At this time, the SGI Origin 3900 is divided into 4 separate nodes. The system has a Symmetric MultiProcessing (SMP) architecture with 512 CPUs, 512 GB memory per node and a total of 40TB of disk space.

The SGI Origin 3900 is front-ended by an SGI Origin 3800 for interactive use and batch submittal. The hardware specifications for this front-end system are slightly different than the batch systems. The front-end SGI Origin 3800 is a Shared Memory system with 12 CPUs per node. Each CPU is a 600 MHz R12000 MIPS processor and has a 32 KB primary data cache, a 32 KB primary instruction cache, and 8 MB on-board cache.

2.2 SGI Origin 3900 File System Overview

Diskspace is subdivided into several areas:

- System space (i.e., /usr, /opt.)
- /workspace
- /scivis

2.2.1 SGI Origin 3900 /workspace

/workspace is a filesystem local to each machine that users are required to run their jobs in.

I/O on the /workspace filesystems is quicker than I/O on a file system mounted from the network (such as your MSRC home directory or archive directory). /workspace is intended for the **temporary** storage of data files needed for your application. This includes (but is not limited to) grid files, restart files, input files, and output files. /workspace is to be used rather than the /tmp and /usr/tmp directory areas.

Each user is given a directory named /workspace/username where username is the user's login name. This directory is created for the user at login via the global .cshrc file, if necessary. The global .cshrc file also creates the environment variable \$WRK which is set to /workspace/username. The user may use this environment variable as needed.

Workspace on the SGI Origin 3900 totals 26 TB in size and utilizes a Redundant Array of Independent Disks (RAID) and uses Clustered XFS (CXFS) to manage /workspace. All nodes have access to /workspace.

There is no quota on the amount of disk space a user may use in `/workspace`, but a file scrubber is used to automatically remove old files from `/workspace` to prevent it from becoming filled. The current policy for removing files is on the ASC MSRC web page. This policy is subject to change based on periodic reviews.

The `/workspace` file system is NOT backed up. In the event of deletion or catastrophic media failure, files and data structures are lost. It is the user's responsibility to transfer files that need to be saved to a location that allows permanent storage. Two possible locations are the user's `$HOME` directory space on the file server or the user's `$ARC` directory on the archival storage system.

2.2.2 SGI Origin 3900 `/scivis`

`/scivis` is a filesystem local to the SGI Origin 3900 that is cross-mounted to the `svw10` and `svw11` SciVis systems using the same CXFS filesystem that mounts `/workspace`. This was done to ease the burden of transferring files between the SGI Origin 3900 and the SciVis systems for the purpose of visualization.

I/O on the `/scivis` filesystems is quicker than I/O on a file system mounted from the network (such as your MSRC home directory or archive directory). `/scivis` is intended for the **temporary** storage of data files needed for your application or for visualization. This includes (but is not limited to) grid files, restart files, input files, and output files. `/scivis` or `/workspace` is to be used rather than the `/tmp` and `/usr/tmp` directory areas.

`/scivis` on the SGI Origin 3900 totals 6.4 TB in size and utilizes a Redundant Array of Independent Disks (RAID) and uses Clustered XFS (CXFS) to manage `/scivis`. All nodes have access to `/scivis`.

Each user is given a directory named `/scivis/username` where `username` is the user's login name. This directory is created for the user at login via the global `.cshrc` file, if necessary. The user may use this environment variable as needed.

There is no quota on the amount of disk space a user may use in `/scivis`, but a file scrubber is used to automatically remove old files from `/scivis` to prevent it from becoming filled. The current policy for removing files is on the ASC MSRC web page. This policy is subject to change based on periodic reviews.

2.3 SGI IRIX Operating System

The SGI Origin 3900 runs IRIX, a 64-bit, multiuser/multitasking operating system based on components from Berkeley Software Distribution (BSD) versions 4.3 and 4.4 and AT&T UNIX System V, Release 4.0.

2.4 Available Software

Software currently available on the SGI Origin 3900 includes: the MIPSPro FORTRAN 77, FORTRAN 90, C, and C++ compilers; MPI and many third party

software packages. A complete list of software is maintained on the ASC MSRC web page at

<http://www.asc.hpc.mil/software/>

2.5 Development

The SGI Origin 3900 has several tools available to help develop, compile, and analyze programs.

The compilers on the SGI Origin 3900 are optimizing and parallelizing compilers that can generate 64-bit and 32-bit codes. The FORTRAN 90 compiler also supports High Performance Fortran (HPF) directives.

The SGI Origin 3900 has the SGI Scientific Computing Software Library (SCSL). This is a collection of mathematical and scientific libraries including Basic Linear Algebra Subprograms (BLAS) levels 1, 2, and 3; LAPACK; Fast Fourier Transforms (FFTs); and convolutions.

The ProDEV Workshop debugger, *cvd*, is a fully symbolic debugger with a graphical user interface. This debugger offers the following features:

- perform source-level debugging
- attach to running process
- debug programs with shared libraries
- debug multithreaded applications
- debug multiprocessor applications, including programs that fork/exec

While a Graphical User Interface (GUI) can help debugging program issues, running an X-window session over a network can be cumbersome. In this case, command line debuggers would be a better choice. The SGI Origin 3900 has both *dbx* and *gdb* installed.

Documentation for these compilers, libraries, and tools is available online in the man page by executing a man on *f90*, *f77*, *cc*, *cxx*, *scsl*, *cvd*, *dbx*, and *gdb*.

3. Program Development

Program development in the SGI Origin 3900 computing environment is similar to that used in a typical UNIX environment. However, the user must take additional steps to utilize the multiple processors available.

3.1 Parallel Processing

Users may utilize multiple CPUs to execute their programs. The compilers are capable of creating parallel programs through the use of compiler directives and parallel standards such as Message Passing Interface (MPI) and OpenMP.

3.1.1 MPI

The goal of MPI is to develop a widely used standard for writing message-passing programs. As such the interface attempts to establish a practical, portable, efficient, and flexible standard for message passing.

You can compile MPI programs on the SGI Origin 3900 using the following command line options.

For MPI FORTRAN codes:

f77 -o prog prog.f -lmpi

f90 -o prog prog.f -lmpi

For MPI C/C++ codes:

cc -o prog prog.c -lmpi

CC -o prog prog.c -lmpi

Other compiler options are available. Please consult the man pages for the compiler you are using for more information.

More information on MPI can be obtained from:

<http://www.mpi-forum.org>

3.1.2 OpenMP

OpenMP is a specification for a set of compiler directives, library routines, and environment variables that can be used to specify shared memory parallelism in FORTRAN and C/C++ programs.

Creating an OpenMP program is done through OpenMP directives in the source code and by adding the -omp flag to your compile string.

To run an OpenMP program, you must first tell the program how many threads (processors) to use. This is achieved through the OMP_NUM_THREADS environment variable. To set this variable, use the following command in *csh*:

setenv OMP_NUM_THREADS x

where *x* is the number of CPUs you wish to run on.

For more information on OpenMP and its directives, please see the following page:

<http://www.openmp.org>

3.2 FORTRAN Programming

The default FORTRAN compiler on the SGI Origin 3900 is the SGI MIPSPro FORTRAN compiler. The FORTRAN compiler commands are *f77* and *f90*. These optimizing and parallelizing compilers can generate 64-bit and 32-bit code. The FORTRAN 77 compiler is fully compliant with the ANSI X3.9-1978 and FIPS PUB 69-1 standards defined for FORTRAN 77. The FORTRAN 90 fully complies with the ANSI X3.198-1992 and ISO/IEC 1539:1991(E) standards definition for FORTRAN 90.

Compiling a FORTRAN program on the SGI Origin 3900 is similar to compiling a program on a typical UNIX system.

f90 -o prog prog.f

This command creates an executable called *prog*. The program is run by typing the program name at the system prompt.

./prog

Further optimization is available through the use of compiler flags and compiler directives. Please check the *f77* and *f90* man pages for more details.

3.3 C/C++ Programming

The SGI MIPSPro C and C++ compilers are available on the SGI Origin 3900. These compilers are capable of optimizing and parallelizing code. The C compiler fully complies with the ANSI X3.159-1989 and ISO/IEC 9899:1990 standards definition for C. The C++ compiler fully complies with the ISO/IEC 14882:1998 standard definition for C++.

Compiling a C program on the SGI Origin 3900 is similar to compiling a C program on a typical UNIX system.

cc -o prog prog.c

Compiling a C++ program is also similar to other systems.

CC -o prog prog.C

These commands will create an executable program in a file called *prog*. The program is executed by entering

./prog

Further optimization is available through the use of compiler flags and compiler directives. Please consult the *cc* or *cxx* man pages for more details.

3.4 Libraries

3.4.1 Math and Science Libraries

The SGI Origin 3900 has SCSL, a collection of mathematical and scientific libraries including BLAS levels 1, 2, and 3; LAPACK, a collection of solvers for dense linear algebra problems, including linear equations, linear least squares problems, eigenvalue problems, and singular value decomposition problems; Fast Fourier Transforms (FFTs); and convolutions. Both single-threaded and multi-threaded routines are available and select routines have been highly optimized to greatly improve performance. Users should use these library routines whenever possible.

This library is not automatically included in the link path. The user must specify the library when linking as in the following examples.

```
f90 -o prog prog.f -lscs  
cc -o prog prog.c -lscs
```

Please consult the *scsl* man page for more details on the routines that are available.

4. Running Jobs on the SGI Origin 3900

4.1 Interactive Use

Interactive use is allowed, particularly for program development, including debugging and performance improvement, job preparation, job submission, and the preprocessing and postprocessing of data. However, only one node on the SGI Origin 3900 system is available for interactive use and interactive jobs are limited to 4 CPUs with 15 minutes of CPU time per process. Jobs with larger resource requirements must be submitted to the batch queues.

4.2 Batch Use

The SGI Origin 3900 uses the Load Sharing Facility (LSF). LSF is a networked subsystem for submitting, monitoring, and controlling a work load of batch jobs on one or more systems. It provides services to monitor queue activity and to delete queued or running jobs. In the event of an orderly system shutdown, LSF jobs will be rerun from the beginning of the job (unless they are specifically marked not to be rerun). More information about LSF is available at

<http://www.platform.com/products/HPC/>

To allow users to run longer in the queues, the SGI Origin 3900 utilizes a 2 week (336 hour) queue, however, to keep the queue structure fair to all users, several restrictions have been put into place:

- The primary resource to schedule jobs by is the cpu hour (CPH). This quantity is equal to the wall time requested multiplied by the number of cpus and cannot exceed a value of 25,000 per user.
 $(ncpus * walltime) = CPH$

Table 1: Example CPH Values

# of Jobs	# of CPUs	Walltime	Total CPH
1	508	48	24,384
2	256	48	24,576
25	8	96	19,200

- A user can use at most 508 cpus.
- The maximum wall time of any queue shall not exceed 336 hours.
- A user can have no more than 25 jobs running at one time no matter what queues they are in.
- A background job cannot start if there is a foreground job in any queue.

The list of queues and the upper limits of job resources for these queues are available on the web at

http://www.asc.hpc.mil/overall/policy_procedure/policies/batchqueue.php

These limits are subject to change based on periodic review of system utilization and system configuration.

TIP: Because of CPH, it is not recommended that users accept the queue default walltime. If your job requires less time to run than the queue default, requesting the smaller of the two will result in a lower CPH and will allow you to **run more jobs** and **reduce your queue wait time**.

Example: User Joe submits a 4 CPU job using the queue default walltime of 336 hours. This results in a CPH of 1344. User Bob submits the same 4 CPU job, but only requests the 120 hours he needs to finish the job. His CPH would be 480. This allows user Bob to submit twice (2.8 to be exact) as many jobs as Joe in the same amount of CPH. If user Bob can alter his walltime and fit his jobs within 110 hours, he can submit 3 times the number of jobs as Joe. This also has a domino effect on the queue structure, resulting in faster throughput throughout the entire system.

4.2.1 Queues Available on the SGI Origin 3900

The ASC MSRC SGI Origin 3900 has three user queues: debug, regular, and background. These queues are available 24 hours a day, 7 days a week.

The debug queue accepts jobs that require up to 32 CPUs, 1 hour of CPU time, and 32 GB of memory. This queue is intended for short runs.

The regular queue is available for production work. Jobs that are submitted without any queue specified will go to the regular queue. The regular queue is divided into subqueues, but users do not submit jobs directly to these subqueues. Rather, the user specifies the number of CPUs, the CPU time, and memory requirements using the `bsub` options below. The job is then routed to the appropriate subqueue.

The background queue is also available for production work. Jobs run in the background queue are not charged against a user's allocation. However, jobs in the background queue are only started when utilization of the machine is low and never when foreground jobs are waiting.

The list of queues and the upper limits of job resources for these queues are available on the web at

http://www.asc.hpc.mil/overall/policy_procedure/policies/batchqueue.php

These limits are subject to change based on periodic review of system utilization and system configuration.

4.2.2 Preparing Jobs

Before a user submits a job, they should prepare a job script. A job script is a UNIX shell script that contains all the commands the user will execute during the job. LSF will place the error and output files in the directory the job was submitted from, so scripts must be written with this in mind. Here is a sample job script.

```

#
#Change to WORK_DIR directory and copy input file.
#
cd $WORK_DIR
archive get -C {directory in $ARC} {filename}
#
#Run the analysis.
#
jstat -l
{My Program}
jstat -l
#
#Archive output and remove $WORK_DIR
#
tar cvf ../{output filename}.tar .
archive put -C {directory in $ARC} ../{output filename}.tar
rm -rf $WORK_DIR
#
#Exit the script.
#
exit

```

This script copies an input file and a program to the user's \$WORK_DIR directory. The \$WORK_DIR directory is a directory created by LSF for users to run their batch jobs. This directory has certain protections from the /workspace scrubber, as long as the job is running, plus five days after it finishes, this directory will not be removed. Then the script changes to the \$WORK_DIR directory, runs the program, copies two output files to permanent storage (one to the \$HOME directory, one to the archival storage system \$ARC), and then deletes the remaining files, the \$WORK_DIR directory and exits.

NOTE: \$WORK_DIR only exists in LSF, you will not be able to change to that directory using the variable \$WORK_DIR.

4.2.3 Submitting Jobs

Once a job script is prepared, the bsub command is used to submit the script to LSF. The command has the following syntax:

bsub < script

Some important LSF options used on the SGI Origin 3900 are as follows (type man bsub for a complete list of options available):

- q *queue* Specifies the name of the queue to which the job will be submitted. For a list of allowable queues, please see:
http://www.asc.hpc.mil/overall/policy_procedure/policies/batchqueue.php
- n *n* Specifies the number of CPUs the job will use.

- W *hh:mm* Specifies the time limit for the job in walltime.
The time should be specified in the hh:mm format
(e.g., 15:00).
*This is a required field, there is no default.
- o *outfilename* Standard output (stdout) for the job is written to
outfilename.
*If you do not specify an output filename, LSF
emails the output to your ASC email account.
- e *errfilename* Standard error (stderr) for the job is written to *err-*
filename. The default name is *jobname.ennn*
where *nnn* is the LSF identifier.
*If you do not specify an error filename, LSF
emails the error information to your ASC email
account.
- J *jobname* Specifies the name of the job.
- P *account* Specifies the account number to charge to.

When a job is submitted to LSF, a unique identifier is assigned to the job by the batch system similar to below:

```
2079.hpc11-0.asc.hpc.mil
```

This identifier is needed when deleting a job.

Options of `bsub` commands are specified within the script file itself. The options are specified using syntax similar to PBS, but each line that contains an option must begin with the `#BSUB` string. Options that are specified within the script file must precede the first executable shell command of the file as in the following example.

```
#!/bin/csh
#BSUB -q regular
#BSUB -n 1
#BSUB -W 168:00
#BSUB -J test
#BSUB -o test.out
#BSUB -e test.out
#BSUB -P WP+WPASC00000000**
```

*Bold headers indicate required fields

**This is an example number. To find your account
number, check your `$ACCOUNT` variable using

```
echo $ACCOUNT
```

More sample batch scripts can be found at the following URL:

```
http://www.asc.hpc.mil/customer/userdocs/samples/  
samplebatch.php
```

4.2.4 Monitoring Jobs

The *bjobs* command is used to report the status of the batch jobs that are currently queued or running. Type `man bjobs` for information about *bjobs* and the options that are available.

The *bjobs* command lists all jobs that are running and queued.

bjobs -u all

<u>JOBID</u>	<u>USER</u>	<u>STAT</u>	<u>QUEUE</u>	<u>FROM_HOST</u>	<u>EXEC_HOST</u>	<u>JOBNAME</u>	<u>SUBMIT_TIME</u>
3373	user1	PEND	default	hpc11-0		test	Feb 5 15:22
3971	user2	RUN	default	hpc11-0	hpc11-3	test	Feb 5 15:22

Here is an explanation of the fields in the *bjobs* output.

Table 2: Fields from *bjobs*

<u>Item</u>	<u>Meaning</u>
JOBID	A unique identifier that consists of the original request number and the machine from which the request was submitted. Format is <i>nnn</i> , where <i>nnn</i> is an integer.
USER	Username of person submitting the job.
STAT	Job status. “RUN” indicates the job is running; “PEND” indicates the job is queued.
QUEUE	Name of the queue where the job is waiting or executing.
FROM_HOST	Cluster domain from where the job was submitted from.
EXEC_HOST	Cluster domain where the job is running.
JOBNAME	Name of the job. This is either the name of the script file submitted to LSF or the name chosen with the -J flag.
SUBMIT_TIME	The date and time the jobs was submitted on.

4.2.5 Deleting Jobs

In LSF, queued or running jobs are removed using the *bkill* command. The syntax is

bkill request-id

where *request-id* is the LSF identifier number.

Example:

bjobs

<u>JOBID</u>	<u>USER</u>	<u>STAT</u>	<u>QUEUE</u>	<u>FROM_HOST</u>	<u>EXEC_HOST</u>	<u>JOBNAME</u>	<u>SUBMIT_TIME</u>
3373	user	PEND	default	hpc11-a	hpc11-b	test	Feb 5 15:22

bkill 3373

5. Customer Service

5.1 Customer Service Center

For customer assistance, call the ASC MSRC Service Center at 1-888-MSRC-ASC (1-888-677-2272), (937) 255-0194, or DSN 785-0194, or send e-mail with a description of the problem to msrchelp@asc.hpc.mil. The support analysts will help with anything related to ASC MSRC: third party software, UNIX, the different ASC MSRC computers, etc. If you have any questions about the ASC MSRC, contact the Service Center first. If your problem or question is beyond the scope of their expertise, they will refer you to the appropriate resource.

5.2 ASC MSRC Support

In-depth technical inquiries and problems are forwarded to the ASC MSRC Customer Assistance and Technology Center (CATC), which pursues such inquiries and problems through resolution as rapidly as possible. The ASC MSRC CATC will attempt to determine the nature of the problem, then identify and coordinate whatever resources are needed to resolve the problem.

The ASC MSRC also offers training classes, which provide an introduction to UNIX and the ASC MSRC. Intermediate and advanced classes on selected topics are also periodically announced on the Programming Environment and Training (PET) section of the ASC MSRC homepage. Topics for such classes may be requested through the Customer Service Center.

The ASC MSRC CATC is ready to support in an advisory capacity any engineer or scientist who is (or potentially is) an ASC MSRC user.

5.3 ASC MSRC Website

The ASC MSRC website is the best source for current ASC MSRC information. To access the ASC MSRC website simply access this URL: <http://www.asc.hpc.mil>.

Some of the topics found on the website include:

APPLICATIONS

Short and long descriptions of current ASC MSRC applications

<http://www.asc.hpc.mil/software/>

SYSTEMS

Information on ASC MSRC servers and Archival Storage

<http://www.asc.hpc.mil/hardware/>

CUSTOMER SERVICE

Available Customer Services

<http://www.asc.hpc.mil/customer/>

ONLINE DOCUMENTATION

Listings of the ASC MSRC User Guides are available for viewing. Instructions are given on obtaining postscript versions.

<http://www.asc.hpc.mil/customer/userdocs/>

VISUALIZATION LAB INFORMATION

Current status and other information about the Visualization Lab.

<http://www.asc.hpc.mil/sciviz/>

TRAINING

Current course offerings and schedule

<http://www.asc.hpc.mil/education/training/>

FREQUENTLY ASKED QUESTIONS

Submit questions and read about various topics (such as “Customizing Your Environment”)

<http://www.asc.hpc.mil>

POLICIES AND PROCEDURES

The latest policies regarding usage of the ASC MSRC resources.

http://www.asc.hpc.mil/overall/policy_procedure/

Appendix A. SGI Origin 3900 Usage Hints

The following are tips and hints for the effective use of the SGI Origin 3900.

A.1 Runtime Considerations

A.1.1 Batch use is recommended

The SGI Origin 3900 system allocates more resources to batch jobs than for interactive use. Users will obtain the best throughput for long running or large memory jobs by submitting jobs to the batch queues.

A.1.2 Request only the time and memory needed

When submitting a job, choose the smallest queue that accommodates the job's time and memory requirements. Jobs that request significantly more resources than are actually needed can result in longer wait times and inefficient use of the machine.

A.2 Files and Filespace

A.2.1 File Management in `/workspace` and `/scivis`

A file scrubber is used to automatically remove old files from `/workspace` and `/scivis` to prevent them from becoming filled. The policy for removing files from these filesystems is available on our website. However, users are encouraged to remove files from `/workspace` and `/scivis` when they are no longer needed. This will minimize the overhead needed to enforce this policy.

The `/workspace` and `/scivis` filesystems are not backed up. It is the user's responsibility to transfer files that need to be saved to a location that allows permanent storage. Two possibilities are the user's `$HOME` directory space on the file server or the user's `$ARC` directory on the archival storage system.

A.2.2 Archival Storage

The directories of the ASC MSRC Archival Storage system are NFS mounted on the SGI Origin 3900. However, copying or moving (i.e., `cp` or `mv`) large files to or from an NFS file system is very slow. `ftp` and `rcp` afford the quickest means to transfer files to and from the archival storage system.

To open an `ftp` session from the SGI Origin 3900 to the archival storage system, type the following command.

```
ftp $msas
```

The `archive` command is more convenient to use than *ftp* within a script. The first example below copies a file from the archive storage system to the local system. The second example copies a file from the local system to the archive storage system.

archive get archive_filename local_filename

archive put local_filename

archive_filename and **local_filename** can be the same. For more details about the Archival Storage system, see the *Archival Storage User's Guide*, located at:

<http://www.asc.hpc.mil/customer/userdocs/guides.htm>

A.2.3 Keep I/O local to the system

The /workspace filesystems are local to the SGI Origin 3900 via a RAID 5 storage. Although \$HOME is NFS-mounted internally to the compute nodes, I/O access from /workspace will be faster than from the HAFS.

Here is a sample script that copies two input files (one from the \$HOME directory, one from the archival storage system) and a program to the user's /workspace directory, changes to the /workspace directory, runs the program, copies two output files to permanent storage (one to the \$HOME directory, one to the archival storage system), and then deletes the remaining files.

```
cd $WRK
cp $HOME/small.input $WRK
archive get big.input $WRK
archive get prog $WRK
prog
archive put big.output
mv small.output $HOME
rm small.input big.input prog big.output
```

A.3 Helpful SGI Origin 3900-related Websites

A.3.1 SGI

<http://www.sgi.com/>